

Secure Industrial Intelligence Across the Value Chain

Practical strategies for manufacturers and industrial
leaders to secure data and operations



Introduction

Secure industrial intelligence across the value chain


Industrial transformation is moving from isolated pilots into production environments, where it helps design products, run factories, coordinate supply chains, and support frontline teams. This shift creates opportunity but also changes what must be protected. Industrial organizations need security that connects identities, data, applications, devices, infrastructure, and AI without disrupting critical physical processes.

This guide explores how industrial leaders can advance frontier transformation by embedding intelligence across the value chain. Each chapter examines a priority business scenario, and the security decisions needed to scale AI responsibly across IT and operational technology.


Manufacturing value chain



Engineering



Production



Logistics



Frontline operations

Chapter 1

Understand new industrial risk within the landscape

Industrial leaders already use intelligence across engineering, operations, maintenance, planning, quality, and workforce support. As systems connect, decisions and actions move across the value chain faster. Security is no longer simply a cost center or separate technical workstream; it enables AI transformation without unmanaged risk. That requires security built deeply into the platform and integrated with specialized partner capabilities to support resilience, regulatory readiness, and growth.

As interoperability grows across IT, OT, suppliers, and cloud services, risk can rise—making secure, trusted connections essential. A compromised identity, unmanaged supplier account, or exposed remote-access pathway can lead from enterprise systems into production. Aging assets, long equipment lifecycles, and limited maintenance windows also make many traditional IT security practices difficult on the factory floor.

The consequences of industrial risk extend beyond data loss—a poorly timed scan, update, or automated response can interrupt production or affect physical safety. Security controls must therefore reflect operational priorities, including continuous availability, safe recovery, clear accountability, and support for modern and legacy systems.

AI adds complexity because models and agents may access operational data, invoke tools, recommend actions, or coordinate workflows. Leaders need visibility into which systems use AI, what data they access, what actions they can take, and how decisions are reviewed and audited.

Regulations and frameworks—including NIS2, IEC 62443, the Cyber Resilience Act, NERC CIP, NIST guidance, and the EU AI Act—increasingly shape investment and architecture decisions.

Disconnected point solutions can add risk through implementation errors and weak seams between siloed systems. Complemented by well-architected partner capabilities, an integrated security and governance spine across identity, data, infrastructure, operations, and AI. Together this reduces fragmentation and supports security, compliance, sovereignty, and transformation.



Questions to bring to leadership discussions

- Which identities and access pathways can affect production?
- Which operational assets are visible to the security team today?
- Where does sensitive industrial data move across sites, clouds, and suppliers?
- Which AI systems can recommend or take operational action?
- Can the organization demonstrate how access, changes, and AI decisions are governed?
- How quickly can IT, OT, security, and operations coordinate during an incident?

Advance security and AI together

Security transformation is most effective when it moves in lockstep with AI transformation. With aligned pacing, both can advance in parallel while risk remains actively managed, acceptably bounded, and proportionate to the business outcome.




Identities
and access



Aging operational
assets



Regulatory
Pressure



AI-enabled
actions

Chapter 2

Build trusted, sovereign, and secure operations

An effective industrial security strategy unifies identity, access, asset health, governance, and response. This reduces exposure while preserving the uptime, safety, and local control manufacturing environments require.

Prioritize the assets and pathways that can change a physical process. Engineering workstations, controllers, safety systems, privileged identities, and remote-access connections need the strongest controls. Verified identity and least-privileged access help ensure that people, devices, workloads, and AI agents access only what they need.

Zero Trust applies the same principles across IT and operational technology, but at different levels of scope. IT controls can be highly granular—applied per identity, application, or data item—while legacy OT assets may need to be secured in bounded groups or zones. The goal is to limit the blast radius of a compromise to an acceptable business and operational cost, measured in time, money, materials, safety, or reputation. Passive discovery, segmentation, brokered access, protected administrative workstations, and compensating controls help manage that risk as plants modernize in phases.

Treat sovereignty as an architecture decision, not simply a hosting preference. Operational outcomes like AI integrated with factory production determine which workloads and data must remain local. An adaptive cloud architecture supports local execution and centralized governance without requiring organizations to predict every future need, providing flexibility across sites, clouds, and edge environments.

Govern AI agents through the same security model: give each a distinct identity, narrowly scoped permissions, approved action boundaries, and an activity trail. Verify access explicitly, treat prompts and responses as potential risks, and monitor agent behavior alongside human and workload activity.



Six capabilities to connect security across IT, OT, and AI

1. Strong identity

Verify users, administrators, workloads, devices, and agents before granting access.

2. Least-privileged access

Replace persistent access with time-bound, task-specific permissions.

3. Device and exposure health

Build an accurate inventory and prioritize the assets that can affect production.

4. Governed change

Evaluate updates and configuration changes against operational safety and uptime requirements, keeping the human in the loop.

5. Coordinated monitoring and response

Connect IT and OT signals so teams can investigate and respond with shared context.

6. Auditable policy enforcement

Apply centrally governed controls consistently across change cycles, with evidence that shows where policy is or isn't working as intended.



Ford builds trust across global operations with Microsoft Defender

Ford deployed Microsoft solutions, including Microsoft Defender, Microsoft Sentinel, and Microsoft Purview, to increase visibility, automate responses, and strengthen data governance across its hybrid environment.



The Microsoft security stack is more than technology. It contributes to Ford's business in moving faster against cyberthreats and building a more secure future.

Weston Maggetti
Platform Manager, Ford Motor Company




[Read the story](#)

Protect what can change the process

Apply the strongest identity, access, monitoring, and governance controls to the systems and pathways that can directly affect physical operations.

Zero Trust operating model



Chapter 3

Protect agents across R&D and digital engineering

Engineering decisions shape cost, quality, manufacturability, compliance, and time to market. AI agents can review requirements, compare designs, organize simulation results, identify constraints, and reuse institutional knowledge. To scale that value, organizations must protect the intellectual property and decision history behind the digital thread.

Engineering agents rarely need broad access. Provide only the design files, simulation outputs, requirements, and industrial context required for the task. Classification and lineage show where data came from, how it changed, and whether it is approved for use.

Secure-by-design practices must extend to connected products. Engineering teams should address identity, update mechanisms, certificate-backed credentials, exposed services, telemetry, and support lifecycles during design. This reduces costly retrofits and supports regulatory readiness.

Human review remains essential. AI can surface alternatives and constraints, but engineers should retain authority over decisions affecting safety, compliance, performance, and intellectual property. Recording the recommendation, evidence, reviewer, and final decision creates a traceable record for future work.

Sovereignty also applies to engineering data. Decide where simulations run, where results are stored, and which jurisdictions, partners, and agents may access them as part of the architecture from the outset

Design principles for secure engineering agents

- Classify engineering and product data before connecting it to AI.
- Limit access by project, product, geography, and role.
- Separate approved design records from experimental content.
- Protect proprietary prompts, models, and engineering knowledge.
- Record recommendations, evidence, approvals, and changes.
- Validate connected-product identity and update mechanisms.
- Keep sensitive compute and data within approved boundaries.



HARTING reduces design time from weeks to minutes with Siemens and Azure AI

HARTING uses Azure OpenAI with Siemens NX to support generative engineering for industrial connectors. The AI-guided experience translates natural-language requirements into product configurations, reducing configuration time by 95%—from 15–20 minutes to about one minute—and creating custom prototype concepts in minutes instead of weeks.



Having the right partners is crucial, and Microsoft is one of them. Together with Siemens, it's perfect.

Norbert Gemmeke

Managing Director, Business Unit Electric, HARTING



[Read the story](#)

Protect the intelligence behind the product

Give engineering agents enough context to improve decisions—not unrestricted access to the full digital thread.



Requirements
and design



Simulation



Production



Field
feedback



Approval

Chapter 4

Govern physical AI and agents on the factory floor

Physical AI is intelligence operating in the physical world. It perceives through sensors and connected assets, reasons over operational data, and acts through machines and industrial environments. In manufacturing, it can inspect parts, adapt a production cell, coordinate material movement, or optimize a batch or continuous process within defined limits. When intelligence moves from recommendation to action, both value and risk become physical.

The data foundation matters as much as the model. Industrial information spans historians, manufacturing execution systems, supervisory control and data acquisition systems, enterprise systems, maintenance records, and engineering documentation. Classification, lineage, quality, and access policies help agents act on trusted information without exposing sensitive operational data.

Every agent that influences production needs a defined purpose. Leaders should know what data it uses, which tools it can invoke, which decisions require human approval, and how to stop or reverse an action. Design these controls into the solution rather than adding them after deployment. Edge deployment keeps latency-sensitive workloads near the line while cloud services provide centralized management and analytics. Local resilience depends on consistent policy, health monitoring, and visibility into changes across software, models, and devices.

Industrial organizations should scale physical AI through bounded scenarios with clear value, reliable data, defined safety requirements, and measurable outcomes. Keep people in control of high-impact decisions until evidence supports greater automation.

Checklist for a controlled physical AI deployment

- Define the operational outcome and acceptable action boundary.
- Assign a distinct identity to each agent or automated workload.
- Document the data sources, tools, and systems the agent can access.
- Require human approval for safety-critical or irreversible actions.
- Capture only the prompts, recommendations, approvals, actions, and outcomes needed to review important decisions and actions, in accordance with privacy, data-minimization, and retention policies.
- Test failure modes, fallback procedures, and manual overrides.
- Monitor model, data, device, and infrastructure health continuously.



Toyota Industries innovates its paint shop processes with Azure industrial AI

Toyota Industries uses Azure industrial AI and connected factory data to identify the drivers of paint defects and accelerate frontline decisions. Its pilot reduced defects by 25%, cut analysis cycles from five days to under four hours, and created four times more opportunities to resolve quality issues.



Seeing real operators use data effectively helped us understand that this was not just a technical solution, but one that could change how people work and make decisions.

Takuya Hatakeyama
Project Lead, Toyota Industries



[Read the story](#)

Automation needs boundaries

Define what an agent may observe, recommend, approve, and execute before it begins influencing the physical process.



Chapter 5

Give workers secure, trusted guidance

Frontline teams need fast access to procedures, diagnostics, safety information, and expert guidance. AI can surface this knowledge in the flow of work, but the experience must account for shared devices, rotating shifts, limited connectivity, and varying authorization.

Identity is the starting point. Shared accounts obscure who accessed sensitive information or approved an action. Individual sign-in, device controls, and role-based access help workers receive appropriate guidance while protecting operational systems and company knowledge.

Content governance determines what an assistant may surface. Safety instructions, operating procedures, engineering documents, maintenance records, and production data have different owners and sensitivity levels. Classifying information before connecting it to an assistant reduces oversharing and preserves a reliable source of truth.

Guided troubleshooting can extend scarce expertise across shifts and locations by helping workers find documentation, understand likely causes, and follow approved steps. For higher-risk tasks, the assistant should escalate to a qualified person rather than act beyond its scope.

The same principle applies to security operations. Unified signals and AI-assisted investigation help small teams cover large environments. Automation should reduce noise and add context while people retain control over actions that could affect production.

Build trust into frontline assistance

- Use individual identities instead of shared accounts.
- Apply access based on role, location, shift, and device health.
- Curate approved knowledge sources and assign content owners.
- Show the source and status of critical guidance.
- Escalate safety-sensitive or uncertain situations to a person.
- Capture important recommendations and worker confirmations.
- Design offline and recovery options for locations with limited connectivity.

Schaeffler unlocks factory insights with Microsoft Cloud for Manufacturing

Schaeffler uses Microsoft Fabric and an Azure AI agent to give factory workers natural-language access to operational data, helping them identify downtime causes faster and act on near-real-time insights across quality, yield, and uptime.



I'm very excited about the agent in Azure AI. It has a lot of potential in the future because it delivers a streamlined end-to-end user experience within the Microsoft platform.

Timo Spiekermann
Product Owner, Data Platform, Schaeffler



[Read the story](#)

Make knowledge available without making it uncontrolled

Give workers fast access to approved guidance while keeping sensitive data, actions, and systems within clear boundaries.



Chapter 6

Secure the agentic supply chain

Supply chains rely on collaboration among suppliers, logistics providers, contract manufacturers, customers, and internal teams. These connections create value, but persistent vendor accounts, shared credentials, and broad network access also create attack paths.

Modern access should be specific to the person, organization, system, task, and time period. A supplier may update shipment status without accessing production systems; a maintenance partner may receive temporary access to one asset at one site. These boundaries reduce risk without slowing collaboration.

AI agents add new cross-organizational actions, from recommending suppliers and adjusting inventory to proposing production changes or initiating transactions. Each action needs an accountable identity, defined authority, and reviewable record. High-impact commitments should require approval and policy checks.

Manage supply chain resilience and cybersecurity together. A cyber incident can interrupt production, block logistics, delay orders, or undermine trust as quickly as a physical disruption. Joint scenarios can test how teams, systems, and agents respond when a supplier is unavailable, data is compromised, or a critical connection must be isolated.

A connected data foundation supports faster decisions without exposing all information to every participant. Classification, contractual boundaries, regional requirements, and business need should determine what users and agents can see.

Questions for suppliers and partners

- Does every external user have an individual identity?
- Is access limited to a site, system, asset, and defined task?
- Can access expire automatically?
- Are supplier and agent actions logged and reviewable?
- Can the organization isolate a connection without stopping wider operations?
- Are data-sharing boundaries aligned to contracts and regional requirements?
- Have teams rehearsed cyber and operational disruption scenarios together?

Dow reimagines productivity and supply chain efficiency with Microsoft 365 Copilot

Dow adopted Microsoft 365 Copilot to empower employees with AI-driven insights and automate tasks across departments. Copilot supports logistics by reviewing freight rates, flagging billing discrepancies, and optimizing workflows for efficiency.



By optimizing this process, we anticipate Copilot will save us millions of dollars on shipping operations in the first year.

Tami Sowle

Logistics Operational Excellence Manager, Dow



[Read the story](#)

Treat every agent action as a business commitment

When an agent orders, releases, schedules, or negotiates across company boundaries, identity and approval matter as much as integration.



Chapter 7

Start with one site and one priority risk

Industrial security transformation need not begin with every plant, platform, and use case. A focused start helps leaders prove value, limit disruption, and create a repeatable model.

Choose one site and scenario with clear business impact, such as third-party access, AI-assisted maintenance, frontline guidance, engineering data, supply chain planning, or coordinated IT and OT incident response. Define the operational outcome and risk boundary before selecting technology.

Bring security, operations, engineering, IT, legal, compliance, and business owners together early to agree on critical assets, data, identities, actions, and recovery requirements. Shared accountability keeps security and governance from becoming late-stage blockers.

Measure progress in business terms: reduced standing access, better asset visibility, faster investigation, fewer escalations, shorter engineering cycles, lower troubleshooting time, improved uptime, or stronger audit readiness. Pair each outcome with evidence that the control works as designed.

Use the first deployment to create a reusable pattern. Document the architecture, roles, policies, partner responsibilities, approval gates, and lessons learned. The goal is a governed operating model that can expand across sites and use cases.

A practical starting sequence

1. Select one site and one high-value operational scenario.
2. Map the identities, assets, data, connections, and actions involved.
3. Identify the systems that can change the physical process.
4. Define local-control, sovereignty, safety, and compliance requirements.
5. Apply identity, least privilege, segmentation, and monitoring.
6. Establish human approval and audit requirements for AI actions.
7. Test response, fallback, and recovery procedures.
8. Measure the business and risk outcomes.
9. Turn the design into a repeatable deployment pattern.

Your next step

Bring security and operations leaders together for an executive briefing on one plant and one priority operational-risk scenario. Clarify the business outcome, identify essential controls, and define a practical path from assessment to deployment.

Modernize with confidence

Connect security, sovereignty, resilience, and AI governance from the start so every new capability strengthens the operation instead of adding unmanaged risk.



Drive industrial transformation
with **Microsoft for Manufacturing**